



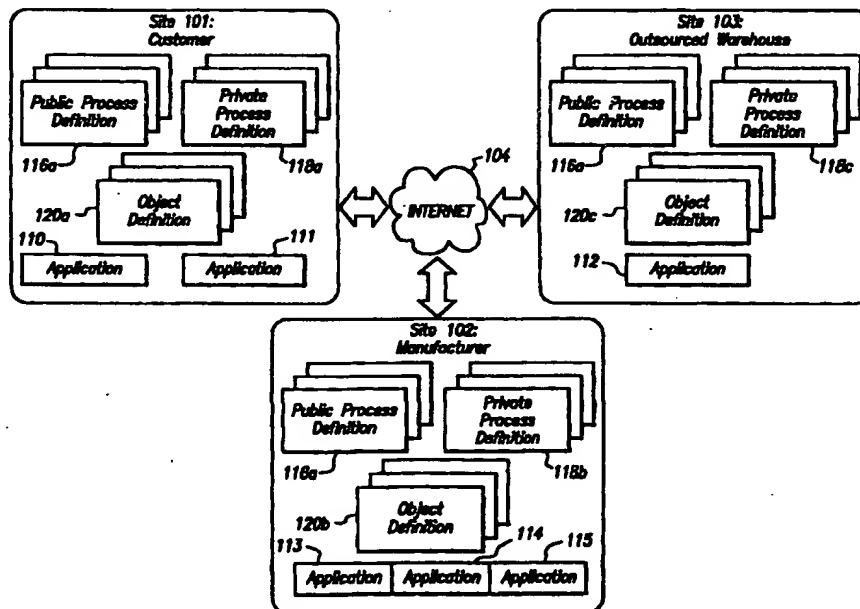
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 15/16		A1	(11) International Publication Number: WO 98/33125
			(43) International Publication Date: 30 July 1998 (30.07.98)
(21) International Application Number: PCT/US98/01403		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, <u>US</u> , VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 23 January 1998 (23.01.98)			
(30) Priority Data: 60/036,385 24 January 1997 (24.01.97) US			
(71) Applicant (for all designated States except US): CROSS-ROUTE SOFTWARE, INC. [US/US]; Suite 190, 555 Twin Dolphin Drive, Redwood City, CA 94065 (US).			
(72) Inventors; and (75) Inventors/Applicants (for US only): OLSEN, Gregory, R. [US/US]; 935 Old County Road #29, Belmont, CA 94002 (US). FROST, Hildreth, Robert [US/US]; 1513 Church Street, San Francisco, CA 94131 (US). THIRUNAVUKKARASU, Chelliah [IN/US]; 2156 Carlmont Drive #1, Belmont, CA 94002 (US). NIBBELINK, Mitchell, W. [US/US]; 299 West Oakwood Boulevard, Redwood City, CA 94061 (US).			
(74) Agents: SUEOKA, Greg, T. et al.; Fenwick & West LLP, Two Palo Alto Square, Palo Alto, CA 94306 (US).		Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	

(54) Title: A SYSTEM AND METHOD FOR CREATING, EXECUTING AND MAINTAINING CROSS-ENTERPRISE PROCESSES

(57) Abstract

A system and method for creating, executing, and maintaining shared, automated business processes across distributed organizations comprises capabilities that enable interoperability among heterogeneous information systems. The system includes a plurality of independent communicating subsystems called sites that have a server with common means of representing and executing shared public process definitions and private process definitions. Process execution comprises coordinated inter-site message exchanges that are coupled with controlled sequences of actions that are local to each of the sites. The public process definition or module captures interactions among the independent sites. Interactions include communication events in which one site sends a message of a known type to another site. Each definition specifies a set of valid sequences of communication events among the participating sites. Associated with any public process definition is a set of lower level or private process definitions or modules. The private process definition specifies a set of possible local actions that can be executed at the site when that particular public process node is executed. In the preferred embodiment, the private process definition is defined in terms of constructs such as operating parameters and software application interactions.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav	TM	Turkmenistan
BF	Burkina Faso	GR	Greece		Republic of Macedonia	TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's	NZ	New Zealand		
CM	Cameroon		Republic of Korea	PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

A System and Method for Creating, Executing and Maintaining Cross-Enterprise Processes

BACKGROUND OF THE INVENTION

1. FIELD OF THE INVENTION

The present invention relates generally to information systems, and in particular to a system and methods that enable the coordination of activity among distributed information systems.

2. Description of the Related Art

The recent change in the structure of business organizations, from the independent monolithic entity to multiple interdependent businesses, mirrors a similar evolution in computer systems from the single mainframe to distributed networks of personal computers and workstations. Since computer networks are extremely efficient in communicating information and performing activities between distributed sites, modern networks should be the obvious beneficiaries of this revolution in technology. For example, such routine activities as ordering and confirming purchases could be performed automatically between existing systems by a shared computer network. Numerous shortcomings in the state of the art, however, prevent full exploitation of conventional network technology for inter-enterprise purposes.

The most pressing problems presented by the use of shared networks between business partners are (1) the heterogeneity of the partner computer systems, (2) the heterogeneity of the data used by the partner systems, (3) communication security and reliability between systems, and (4) the legal, organizational and cultural boundaries among partners.

With respect to system heterogeneity, organizations often use combinations of operating systems, middleware systems, and software applications that are incompatible with one another. Widespread middleware deployment is now underway, but interoperation among the leading camps has yet to be fully defined. Especially problematic are differences at the application level, which are fundamental and will continue to challenge implementers for some time.

With respect to the second problem, data heterogeneity, different applications and users of those applications often represent information in different ways or use different kinds of

1 information to accomplish the same task. These gaps can be particularly significant when the
2 applications and their users are distributed among different enterprises. Bridging the associated
3 syntactic and semantic gaps in information can require a mixture of transformation capabilities
4 as well as neutral objects.

5 With respect to the third problem, communication security and reliability, any
6 interaction among the systems of a business network requires the presence of reliable and secure
7 communication pathways between the participants. The concern for security is especially
8 prominent when the Internet is used as a link in the communication pathway, since this medium
9 is susceptible to eavesdropping and other forms of security attacks.

10 With respect to the fourth problem, any attempt to automate business processes between
11 multiple partners must overcome the numerous non-technical barriers associated with
12 management of a project distributed among multiple organizations. These challenges include
13 mismatches between project priority and resource allocation, language barriers, time zone
14 differences and both corporate and governmental regulations. These challenges limit the levels
15 of coordination that are achievable. Any technical solution, therefore, must focus on
16 minimizing the scope and complexity of the mutual commitments required to implement the
17 solution.

18 Currently, there are at least five known methods for extending interdependent processes
19 beyond one computer system to other systems connected by conventional computer networking
20 resources. The first is the manual approach in which users of multiple computer systems
21 communicate information between one another via telephone, fax, or other media. The
22 communicated information is then entered by hand into the respective computer systems. The
23 manual approach may be used to bridge gaps in automation, but is obviously limited in its
24 ability to tightly couple processes among partners in both reliable and efficient manner.

25 The second approach, which arose in the era of home-grown mainframe applications, is
26 known as Electronic Data Interchange (EDI). EDI is a broadly defined term, but most often
27 refers to a particular set of standards, technologies (Value-Added Networks, Direct Dial-ups,
28 mapping software), and practices used for electronic data exchange among companies. In EDI,
29 a collection of business information (e.g., a purchase order) may be exported from one
30 application system, mapped into a neutral format, transmitted to a partner via a VAN (Value-

1 Added Network), mapped by the partner into a format suitable for its application, and imported
2 into the partner application. Alternatively, direct dialing may be substituted for the VAN. EDI,
3 however, is generally batch oriented, requires extensive format customization and does not
4 support processes.

5 The third approach, used when business requirements do not fit the EDI model, is to use
6 a custom system designed and implemented to the users' specifications. This approach is
7 costly, requires a mixture of network programming and system integration tasks, and serves a
8 specific purpose for specific users only. Furthermore, it is inflexible and difficult to modify.

9 Recently, two trends in technology have radically changed the ways in which application
10 systems may be intertwined. As a result, a fourth and fifth approach, as well as an adaptation of
11 EDI, must be added to the three discussed above. The first trend is the rapid expansion of
12 network infrastructure. The most visible component of this infrastructure is the ubiquitous
13 connectivity provided by the Internet. Nearly all organizations are, or soon will be, connected to
14 the Internet. Coupled with this connectivity is an expanding set of middleware technologies and
15 services such as distributed object frameworks and message oriented middleware that facilitate
16 more tractable distributed applications and promise far greater interoperability of software
17 components.

18 The second trend involves advances in the enterprise application systems that are used
19 by companies. Key advances in these systems include development of object interfaces and the
20 development of workflow/process modeling capabilities. Object interfaces provide a more
21 flexible and less taxing method of moving information to and from applications than prior
22 methods such as SQL (Structured Query Language) or file-based interfaces. Several application
23 vendors now provide the ability to design and implement workflow among different application
24 modules. This capability allows companies to more easily focus on their business processes and
25 makes more obvious the need to connect the processes of business partners.

26 The most visible impact of these trends, and the fourth approach to extending business
27 interdependency, is the use of the World Wide Web for business-to-business interactions. In
28 this model, an employee in one business accesses information, such as catalog or shipping
29 information, pertaining to the business applications of another company by using a standard
30 Web browser. This approach, however, is ill-suited to many extended enterprise processes that

1 require dependent interactions among the application systems of different organizations.

2 The fifth approach exploits recent middleware technology which makes possible the
3 creation of high-performance distributed applications that are logically integrated. Though the
4 same technology may be used to provide interoperation among application suites from different
5 vendors and among systems at different businesses, significant challenges limit feasibility.
6 Foremost, employing middleware technology is a programming task that requires significant
7 programming skill and special understanding of security, synchronization, and other network
8 issues. The cost of such an endeavor may be justified for the vendor of a distributed application,
9 but companies wishing to engage in a specific extended enterprise process are unlikely to devote
10 the capital required to build distributed systems from the ground up.

11 Finally, the adaptation of EDI, referred to as Internet-EDI, is actually a number of
12 methods that attempt to move the traditional EDI approach discussed above to an Internet
13 transport medium. These methods are motivated by the desire to reduce high transport costs
14 associated with Value-Added Networks (VAN's). Effectively, these methods diverge very little
15 from traditional EDI. The same message formats, mapping software, and even enveloping
16 constructs are employed. Use of an open network, however, requires additional security,
17 reliability and auditing capabilities that were formerly part of a VAN service. In addition, the
18 use of these additional services in an open network configuration must be supported by software
19 at the endpoints of an information exchange. Internet-EDI, therefore, suffers from key
20 limitations such as a lack of process support, an unwieldy representation formalism, and an
21 integration model that does not mesh with new practices.

22 The approaches above fail to meet the increasing demand to implement between
23 disparate systems complex, automated processes that are both secure and maintainable. Thus, a
24 system and method to plan and control extended business interdependency are needed that (1)
25 focuses specifically on peer-to-peer interactions among existing business application systems,
26 (2) supports secure and reliable communication, (3) minimizes custom software development,
27 (4) has functionality to handle heterogeneous data representation formalisms and (5) has the
28 ability to support complex processes that extend into and out of enterprise applications.

29

SUMMARY OF THE INVENTION

The present invention is a system and methods for creating, executing, and maintaining cross-enterprise processes. Cross-enterprise processes are shared automated business processes or workflows among distributed information systems that include specific provisions for automation of these processes across organizational boundaries and among heterogeneous information systems.

The system is comprised of a plurality of independent communicating subsystems called sites with common capabilities. Each of the sites includes a server with common means of representing and executing shared process definitions. These sites act in concert in the course of executing shared inter-system processes. Process execution comprises coordinated inter-site message exchanges that are coupled with controlled sequences of actions that are local to each of the sites. In addition, each site may include any one of a number of applications programs and operating systems for executing the inter-system processes and internal processes on the server.

Automated inter-system processes are represented in the system of the present invention in a two-level process model. The top level or public process definition/module captures interactions among the independent sites (each typically representing an organization or business unit). Interactions include communication events in which one site, designated in the public process definition by a node, sends a message of a known type to another site. The public process definition, then, is a logical grouping or directed graph of interdependent communication events among a set of sites. Each definition specifies a set of valid sequences of communication events among the participating sites.

Associated with any public process definition is a set of lower level or private process definitions or modules. A separate private process definition is bound to each node in a public process. The private process definition specifies a set of possible local actions that can be executed at the site when that particular public process node is executed. In the preferred embodiment, the private process definition is defined in terms of constructs such as operating parameters and software application interactions specific to the node or site.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram of an extended enterprise including a plurality of sites having public process definitions and private process definitions in accordance with the present invention.

Figure 2 is a graphical representation of a public process definition including node, arcs and connections between them in accordance with the present invention.

Figure 3 is a flow diagram of the process for executing a private process definition in accordance with the present invention.

Figure 4 is a block diagram of a system in accordance with the present invention.

Figure 5 is a flow diagram of a method for distributing a public process definition in accordance with the present invention.

Figure 6 is a flow diagram of a method for the installation of a public process definition in accordance with the present invention.

Figure 7 is a flow diagram of a method for executing an instance of a specific process type in accordance with the present invention.

Figure 8 is a graphical representation of a display device showing a graphical user interface for editing the public process definition.

Figure 9 is a graphical representation of a display device showing a graphical user interface for editing the private process definition.

DESCRIPTION OF THE PREFERRED EMBODIMENT

Referring to Figure 1, a preferred embodiment of an extended enterprise system 100 is shown. The extended enterprise system 100 of the preferred embodiment of the present invention preferably comprises a plurality of sites 101, 102 and 103 installed at distinct organizations that are coupled by a communications network 104. These sites 101-103 form an extended enterprise 100 in which the internal processes of each site 101-103 are coupled with the internal processes of other sites 101-103 via coordinated sequences of information exchanges. For example, sites 101-103 may be business enterprises that comprise three elements of a supply chain: supplier, manufacturer and customer. Those skilled in the art, however, will recognize that the sites 101-103 could for any type of business unit or function,

1 that there may be any number of site, and three sites 101-103 are provided only by way of
2 example.

3 Each of these sites 101-103 represents a zone of control and is comprised by a set of
4 application systems that store information and contain logic for retrieving and modifying that
5 information. Example applications include ERP (Enterprise Resource Planning) application
6 suites, Product Data Management (PDM) systems, logistics applications, and advanced planning
7 systems (APS).

8 Operation of the present invention includes coordinated sequences of actions within each
9 site 101-103 that are linked with coordinated sequences of information exchanges among
10 different sites 101-103. The actions executed within each site 101-103 include primarily the
11 movement of information into and out of the applications associated with sites 101-103. Each
12 exchange of information between sites 101-103 is preceded by a sequence of actions within the
13 sending site and is followed by another sequence of actions within the receiving site.
14 Accordingly, these site-specific sequences of actions serve as the connections that bind a set of
15 information exchanges into a single coordinated sequence of interactions.

16 The possible sequences of local actions and site to site exchanges are specified through a
17 process definition language. This language allows for complex branching and looping logic and
18 can capture constraints that govern the relationships between local action sequences and site to
19 site exchanges. More formally, the process definition language of the preferred embodiment
20 includes node and arc elements that are combined with a specific ordering and logic to create a
21 directed graph (such as depicted in Figure 2 and as will be described below). A single source
22 node 205-225 and a single destination node 205-225 define each arc element. Each node 205-
23 225 includes a set of input arcs and associated logic connecting it to antecedent nodes 205-225
24 and a set of output arcs and associated logic connecting it to consequent nodes 205-225. The
25 relationship between the input arcs for a given node 205-225 is defined by a logical sentence,
26 containing possibly nested conjunctive and disjunctive propositional connectives, in which each
27 arc is represented by a distinct propositional symbol. The output arcs for a given node are
28 related by a separate logical sentence of equivalent form. Node 201 has no input arcs and is
29 referred to as an initial node. Nodes 299 that have no output arcs are referred to as terminal
30 nodes.

1 In the preferred embodiment, a two-level process model is used to represent the
2 collection of site-to-site information exchanges and site-specific sequences of actions. A public
3 process definition or module 116a specifies the relationship between all site-to-site information
4 exchanges. The sequence of possible actions within a single site 101, 102, 103, for a particular
5 node in a public process definition 116a, is specified by a private process definition or module
6 118a, 118b, 118c. Both the public and private process definitions 116a, 118a, 118b, 118c are
7 built on top of the process definition language with special interpretations for node 205-225 and
8 arc elements. In a public process definition 116a, each node element represents a specific site
9 101, 102, 103 and each arc element represents a message with specific information contents that
10 are sent from the site 101, 102, 103, represented by the arc's source node to the site represented
11 by the arc's destination node. The graph for a public process can include only a single initial
12 node. The public process definition 116a, then, is a specification of "who does what when"
13 among a set of sites 101-103 for a particular purpose. Each public process definition 116a
14 specifies a set of valid sequences of communication events among the participating sites 101,
15 102, 103. More specifically, the same public process definition 116a is provided to each site
16 101, 102, 103 having an action in the public process definition 116a. As shown in Figure 1,
17 each site 101, 102, 103 may have one or more public process definitions 116a, one for each
18 inter-site process. In a private process definition 118a, 118b, 118c, node elements represent
19 specific programmatic actions and arc elements specify the order in which these actions are
20 executed. The private process definition 118a, 118b, 118c specifies how sites 101-103 process
21 received messages and construct outgoing messages. Moreover, private process definitions
22 118a, 118b, 118c specify what happens within a node of public process definition 116a. Thus,
23 the private process definitions 118a, 118b, 118c include routines and processes that are tailored
24 to the particular site 101, 102, 103 to which the private process definitions 118a, 118b, 118c is
25 assigned or operates upon. Still more particularly, the private process definitions 118a, 118b,
26 118c are designed for interaction using the operating systems, applications and resources of the
27 site to which it is assigned. Thus, as depicted in Figure 1 each of the private process definitions
28 118a, 118b, 118c is different for each site 101, 102, 103. Nonetheless, when sites 101, 102, 103
29 are similarly configured (e.g., have the same operating systems, applications and resources) the
30 private process definitions 118a, 118b, 118c may be used or shared. Those skilled in the art will
31 recognize that each site 101, 102, 103 may also include a plurality of private process definitions
32 118a, 118b, 118c as depicted in Figure 1. The plurality of private process definitions 118a,

1 118b, 118c may be for one public process definition 116a or different public process definitions
2 116a.

3 The preferred embodiment models the information contained in the messages sent
4 between sites 101, 102 and 103 as objects with restricted structure and behaviors. These objects
5 are data containers whose possible contents are specified by object definitions 120a, 120b, 120c.
6 In the preferred embodiment, an object definition 120a, 120b, 120c takes the form of an XML
7 (Extensible Markup Language) DTD (Document Type Definition). This definition specifies the
8 lexical and grammatical form of all objects of that type. Object definitions 120a, 120b, 120c are
9 referenced by both public and private process definitions.

10 Referring now to Figure 2, an exemplary public process definition 200 is shown. The
11 public process definition 200 is represented graphically as a flow chart. Figure 2 shows a public
12 process definition 200 that specifies a set of possible interactions among sites 101-103. Process
13 200 comprises a set of nodes 205 through 225 and a set of communication events 230 through
14 250. Each node corresponds with a specific site 101-103, and associated with each node 205
15 through 225 of public process definition 200 is a private process definition 118a, 118b, 118c.
16 Figure 3 shows an exemplary private process definition 300 associated with node 210.

17 Each communication event 230-250 that connects one node to another in a public
18 process definition 200 represents the exchange of a message of a known object type. For
19 example, the known objects may be any one of the conventional types of business objects such
20 as a purchase order object, an confirmation messages object, etc. Such objects 120a, 120b, 120c
21 are defined in the object definition 120a, 120b, 120c so that each site 101, 102, 103 may use the
22 object definitions 120a, 120b, 120c as needed to process objects on either the public level or the
23 private level. In other words, public process definition 200 is a logical grouping or a directed
24 graph of interdependent communication events 230-250 among the sites 101-103 shown in
25 Figure 1. This grouping specifies a set of valid sequences of communication events among the
26 participating sites 101, 102, 103. Specifically, public process definition 200 describes at
27 communication event 230 a purchase order sent from site 101 to site 102. This purchase order
28 is generated by the private process of site 101 associated with node 205. Node 210 is a
29 branching node, and two communication events, 235 and 240, are produced at that node.
30 Depending on a specified branch condition, either one of the events occurs or both events occur.
31 The conditions under which these two events trigger are not indicated by public definition 200,

1 since they are known only to the site 102. These conditions are contained in private process
2 definition 118b associated with node 210, and thus site 102. Those skilled in the art will also
3 recognize that the private process definition 118b may be a set of private process definitions that
4 correspond to an instance of a public process definition 116a.

5 The occurrence of events 235 and/or 240 causes the execution of the node(s) that
6 immediately follow and execution proceeds downward on the Figure. Node 225 represents a
7 branch junction node that can wait for one or all events of a set that connect to it. Following
8 execution of node 225 the public process terminates at terminal 299.

9 Associated with each node 205 through 225 of public process definition 200 is a private
10 process definition 118a, 118b, 118c. For example, Figure 3 shows a private process definition
11 300 associated with node 210. Unlike the public process definition 200, the content of private
12 process definition 300 is determined and known solely by the corresponding site, site 102 in this
13 case. Private process definition 300 includes a number of actions 305 through 330 that are
14 controlled according to a specified logic. Possible actions include external business application
15 interactions, script execution, user notification and approval, time delay, output object
16 specification, and sub-process execution. All instances of the private process definition 300
17 have access to the object of type 'Purchase Order' that is contained in communication event
18 230. Any action in private process definition 300 can reference this object. Private process
19 definition 300 is constrained to produce a object either of type 'Acknowledgment' or of type
20 'Purchase Order' corresponding to communication events 235 and 240 respectively.

21 Referring now to Figure 3, one exemplary embodiment for a private process definition
22 300 is shown. Private process definition 300 begins at initiator action 301 that executes after
23 communication event 230 completes. Those skilled in the art will recognize that the process
24 would be similar for a variety of other private process definitions such that once the object is
25 transferred to or received by the site 101, 102, 103, the private definition corresponding to the
26 node following the communication event is automatically started. Private process 300 continues
27 to action 305. Action 305 execution entails posting an information block corresponding to the
28 received purchase order into business application 113. Following completion of action 305
29 execution continues to action 310. Action 310 entails querying business application 114 to
30 determine if the item associated with the Purchase Order of communication event 230 is locally
31 stocked or if the item is outsourced. The result of this query is put into a variable named

1 'OUTSOURCED' in a set of variables associated with private process 300. Action 315 is
2 executed subsequent to action 310. Action 315 entails an IF-THEN-ELSE conditional test on
3 the value of OUTSOURCED inside of a script action. Results of this conditional test determine
4 whether path A or path B is followed subsequent to action 315 completion. Path A execution
5 proceeds to action 320, which entails the construction of a object of type 'Acknowledgment'
6 and its designation as an output of the private process. Path B execution proceeds to action 325,
7 which entails the construction of a object of type 'Purchase Order' and its designation as an
8 output of the private process. Paths A and B terminate in action 330, which entails notification
9 of a designated user via electronic mail of certain status information associated with the running
10 process. This information includes identifying characteristics of the purchase order and results
11 of business application queries. Following execution of action 330, the private process
12 terminates, and control returns to the public process level. The use of such private definitions is
13 particularly advantages because it provides uniform control and regulation of the inter-site
14 processes, while allowing maximum flexibility through the use of private definitions that allows
15 the controller of a particular site to implement the private definition in any number of ways
16 according to parameters, resources, and other constraints for a particular site.

17 Each site of the preferred embodiment includes a combination of components that
18 support the design, implementation and maintenance of public and private processes and the
19 runtime components that support the execution of these processes. Figure 4 shows the preferred
20 configuration of an example site 102.

21 The standard site 102 is comprised by a single server 480 and one or more clients 460,
22 470 that communicate with the server over a network 409. Clients 460 and 470 and server 480
23 run on separate host computers. Clients 460 and 470 contain graphical user interfaces (GUI's)
24 465 and 475 respectively. In addition, server 480 includes or has access to database 410 and
25 applications 420 and 430. In the preferred embodiment, database 410 resides on a host
26 computer that is separate from the one on which server 480 is located. Clients 460 and 470 and
27 server 480 share common representations of relevant information by interacting over network
28 409 according to a specified and conventional communication protocol. Such shared
29 information representations include public and private process definitions, object definitions,
30 process execution histories, as well as information about other sites with which the site interacts.
31 Human users 440 and 450 interact with site 102 via client GUI's 465 and 475 to view, create,

1 edit, and manage the shared information representations delimited above. For example, users
2 440 and 450 are able to view and edit graphical representations of public process 200 and
3 private process 300 on GUI's 465 and 475. Figures 8 and 9 shown a screen shot of the GUIs
4 corresponding to the public and private process definition described above with reference to
5 Figures 2 and 3, respectively.

6 The server 480 of the preferred embodiment is comprised of a middle-tier manager set
7 482, an execution engine 484, a transport manager 486 and adapters 488 and 489. The middle-
8 tier manager set 482 controls the access and flow of information between network 409, engine
9 484, and database 410. In addition, it implements associated application logic, and insures the
10 consistency of information between these elements. With respect to network 409, set 482
11 mediates access to information from concurrently operating clients and other components of
12 server 480.

13 As discussed in detail below, the installation of public and private processes 200 and 300
14 requires the prior approval of user 440 or 450. Once this approval is received, it is entered by
15 the appropriate user into client 460 or 470 via the respective GUI. A local install signal is then
16 relayed over network 409 to server 480. Manager set 482, acting upon the received signal,
17 initiates the installation of the process definitions into engine 484. During installation, the
18 execution engine 484 transforms the process definitions it receives into executable state
19 machines which are saved in database 410. This transformation extracts from the public process
20 definition all nodes connected to arcs involving the target site. The resulting state machine
21 contains all information necessary for a single site to participate in the execution of the original
22 public process. Once the installation is complete, manager set 482 provides engine 484 with
23 any additional information stored in database 410 or received from clients 460 and 470 needed
24 to perform process execution. Persistence of shared data is maintained by communication with
25 database 410.

26 Once the installation of private and public process definitions 200 and 300 is complete,
27 engine 484 controls their execution.. During execution, the execution engine 484 manages two
28 key activities: inbound and outbound communication with other sites via transport manager 486
29 and interactions with applications 420 and 430 via adapters 488 and 489. The engine 484 also
30 manages through manager set 482 several auxiliary activities including the sending and

1 receiving of messages to and from users 440 and 450 and the storage of log information in
2 database 410.

3 During the execution of a public process definition, such as definition 200, transport
4 manager 486 manages communications to and from Internet 104. For example, public process
5 definition 200 anticipates the reception of purchase order 230 and acknowledgment 245 by site
6 102, as well as the sending of acknowledgment 235 and purchase order 240. In this capacity,
7 manager 486 preferably handles retry and acknowledgment logic (based on the properties of the
8 service it is using). Messages are created outside of any specific transport service and
9 communication security is message based. Non-repudiation receipts for both origin and
10 delivery are supported.

11 During execution, adapters 488 and 489 mediate the flow of data between the execution
12 engine 484 and external applications 420 and 430. For example, referring to step 315 of private
13 process definition 300, engine 484 may transmit a request through adapter 488 or 489 that
14 application 420 or 430 determine whether the item in question is outsourced. In turn, the
15 application will respond through the respective adapter. Adapter configuration options for 488
16 and 489 are set by authors of private processes for the associated site. These adapters 488 and
17 489 communicate their acceptable configuration options to the middle-tier manager set 482 at
18 the time of their installation. The configuration interface for adapters 488 and 489 allows a
19 private process to insert data into an external application, retrieve data from an external
20 application or listen for a specific event produced by the external application, where the
21 inserted, retrieved or listened for data is represented by an object definition. Adapters 488 and
22 489 also insure uniform properties of state/consistency management and auditing behavior
23 across the different applications that can be integrated with the system. During process
24 execution, adapters 488 and 489 map the insertion, retrieval and listened for actions specified in
25 a private process into specific interactions with target applications 420 and 430.

26 Operation of a site of the preferred embodiment revolves around the life cycle of a
27 public process definition and the associated private process and object definitions. Shown in
28 figure 5, this cycle begins with the creation of a public process definition and referenced object
29 definitions then continues with the distribution of the public process definition and object

1 definitions, creation of the necessary private process definitions, installation of the process, and
2 ends with process execution.

3 In step 502, the user creates a public process definition. The site at which the public
4 process definition 200 is created is referred to as the authoring site. Creation of the public
5 process definition 200 includes the creation of all object definitions that represent site-to-site
6 messages in the public process definition. In the preferred embodiment, both public process and
7 object definitions are created by users 440 and 450 interacting with manager set 482 through
8 client GUI's 465 and 475. During creation of public process definition 200, for example, user
9 440 specifies the sequence of interactions among all participating sites 101, 102 and 103 and the
10 logic connecting these interactions. In this example, GUI 465 would present definition 200 as a
11 set of icons interconnected by flow indicators that would appear much as the diagram in figure
12 2. Definitions for the purchase order, acknowledgment and rejection objects would also be
13 created by user 440 via GUI 465 if they did not pre-exist.

14 After the public process and necessary objects have been defined, the user proceeds with
15 distribution of the process. In step 504, the authoring site sends over the internet 104, the
16 authored public process definition and referenced object definitions to all sites participating in
17 the public process. Those skilled in the art will recognize that internet 104 may be a intranet on
18 a local area network (LAN), an internet on a wide area network (WAN), or the Internet. In this
19 case, site 102 is the authoring site and sites 101 and 103 are the participant sites. In order for
20 site 102 to send public process definition 200 and associated object definitions over the Internet,
21 the definitions are sent from manager set 482 to transport manager 486 and from there to the
22 transport managers of the participant sites. Upon receipt of public process definition 200 and
23 object definitions by the participant site transport managers, in this case sites 101 and 103, the
24 definitions are passed from the transport manager to the middle-tier managers for persistent
25 storage in the site database. After reviewing received public process definition 200 and object
26 definitions via a client GUI, users at the participant sites 101 and 103 must approve or
27 disapprove the public process definition, said approval or disapproval being sent via the
28 transport managers of the partner sites to the authoring site. While public process definition 200
29 is being reviewed at participant sites 101 and 103, authoring site 102 waits for the approval or
30 disapproval votes to be received by transport manager 486 in step 508. The approval or
31 disapproval by sites 101 and 103 is likely to turn on commercial, rather than technical, concerns.

1 If a partner site finds the commercial arrangements described in definition 200 acceptable, it
2 returns an approval signal to the authoring site, in this case site 102. In step 508, the system
3 tests for universal approval, if either participant site 101 or 103 disapproves of public process
4 200, authoring site 102 will distribute an abort message through transport manager 486 to
5 partner sites 101 and 103, thus reaching step 510. In this case, public process definition 200 is
6 abandoned and sites 101, 102 and 103 may start negotiating for a new public process definition.
7 If the public process definition is universally accepted in step 508, the authoring site 102
8 distributes a commit message to each of the partner sites .

9 After the commit messages have been transmitted by the authoring site and received by
10 the participant sites, both the authoring site and participant sites proceed with creation of the
11 private processes associated with the public process nodes owned by each site. This is
12 represented by step 514 in figure 5. Each site user creates a private process definition for each
13 node of the public process definition associated with the user's site. For example, user 440,
14 creates private process definition 300 for node 210 and an accompanying private process
15 definition for node 220, since nodes 210 and 220 are each associated with site 102. Likewise, a
16 user at site 101 would create private process definitions for nodes 205 and 225, while a user at
17 site 103 would create a definition for node 215.

18 After successfully implementing the necessary private process definitions, each
19 participant site sends a message to the authoring site signally the completion of private process
20 implementation. In step 516, the authoring site gathers private process completion signals from
21 all sites. If any site fails to implement one or more private processes, it will send a failure
22 message to the authoring site. In this case, the implementation process will be aborted (step
23 518), and process definition 200 will be abandoned. Once the authoring site has received
24 messages from all participant sites indicating successful private process implementation, and
25 has successfully implemented its own private processes, the authoring site can begin the
26 installation process (step 520).

27 Process installation (step 520) begins with the authoring site sending installation
28 messages to all participant sites. After receiving the installation message, each participant site
29 locally installs the public process. Figure 6 shows a flow diagram of a process for installing at a
30 single site the private process definitions associated with a public process definition. In step
31 602, the public process definition and associated private process definitions are passed from

1 manager set 482 to execution engine 484. In step 604, the public process definition is compiled
2 to produce a state machine that contains states only for the site in question. For example, the
3 process of compiling public process definition 200 at site 102 will result in states associated
4 with nodes 210 and 220. Recorded in the state machine is a triggering event for each state.
5 Continuing with the example of public process definition 200, site 102 records that event 230, a
6 purchase order from site 101, triggers the state associated with node 210, and that event 245, an
7 acknowledgment from site 103, triggers the state associated with node 220. In step 606, each
8 state of the state machine is bound by a "call" command with an associated private process
9 definition. For example, site 102 will bind private process definition 300 with the state
10 associated with node 210. The result of this binding is that when the state associated with node
11 210 is triggered by a purchase order from site 101, private process definition 300 is called and
12 executed. In step 608, a determination is made as to whether the site in question is the initiator
13 of the public process. If not, as in the example of site 102, the execution engine determines the
14 triggering message to be received and registers it in the transport manager. In the example of
15 site 102, the two triggering messages are purchase order 230 from site 101 and acknowledgment
16 245 from site 103. If the site is the initiator of a public process, the triggering event for the first
17 private process is internal to the site and is registered in step 612 as an event trigger, a scheduled
18 startup, or as a subprocess trigger.

19 After successfully installing the public process, each participant site sends an installation
20 confirmation message back to the authoring site. In step 522, the authoring site collects
21 installation confirmation messages from all participant sites. If any site is unable to install the
22 public process, the process is aborted 524 as described above. Successful installation at the
23 authoring site triggers the transmission of messages to all participants indicating that the public
24 process has been installed at all involved sites. At this point the process is ready for execution
25 (step 526).

26 As discussed above, the execution of a public process is actually performed by the
27 interactive executions of the associated private processes at the partner sites. The execution of
28 an installed public process by a single site is shown in figure 7. Execution begins in step 702
29 with an initiating event that can include the following: receipt of a message from a partner site,
30 an event associated with an application, a scheduled startup, or a subprocess trigger. For
31 example, node 210, of public process definition 200, is triggered by the receipt of a purchase

1 order message from site 101. In step 704, the execution engine at the triggered site creates an
2 instance in the appropriate state machine and sets the machine to the initial state. In step 706,
3 the execution engine fetches the private process associated with the associated public process
4 node. For example, upon the triggering of node 210, private process definition 300 is accessed.
5 In step 708, the execution engine passes appropriate data including the contents of the event 230
6 to the private process and initiates its execution. In step 710, the private process executes and
7 returns data to the execution engine, and in step 712, the execution engine acts upon the basis of
8 the returned data. For example, private process 300 returns to execution engine 484 either an
9 instruction to send an acknowledgment to site 101 or a purchase order to site 103. It is noted
10 that during the execution of the private process in step 710, engine 484 may make use of
11 applications 420 and 430. Execution engine 484 responds accordingly. In step 714, a
12 determination is made by the execution engine as to whether a local terminal state of the public
13 process definition has been reached. This determination is a local determination and is limited
14 to the participation of the site in the public process. For example, the completion of node 220 is
15 a local terminal state for site 102, since it is the final node in process 200 that corresponds with
16 site 102. Likewise, depending on the outcome of the accompanying private process, the
17 completion of node 210 may be a local terminal state for site 102. If local termination is
18 encountered, the public process ends for the site in step 716. Step 716 is not complete until a
19 two-phase commit protocol has been executed among sites 101, 102, and 103 ensuring mutual
20 completion of process 200 execution.

21 If the local terminal state is not encountered in step 714, the transport manager waits for
22 triggering messages from partner sites (step 718). For example, if the result of node 210 is the
23 transmission of a purchase order in event 240, transport manager 486 waits for the
24 acknowledgment of event 245 to trigger the private process associated with node 220. Once the
25 triggering message is received in step 720, the execution engine looks up the associated process
26 state in step 722. The private process then begins execution at step 706 and is performed as
27 described above.

1 WHAT IS CLAIMED IS:

2 1. A method for coordinating a process between a first site and a second site using a
3 public process definition comprised of a first node associated with the first site, a second node
4 associated with the second site, and an arc connected to the first and second nodes, comprising:
5 executing the first node of the public process definition at the first site by executing a
6 first private process definition associated by shared information with the first
7 node; and,
8 upon receiving the message defined by the arc, executing the second node of the public
9 process definition at the second site by executing a second private process
10 definition associated by shared information with the second node.

11
12 2. The method of claim 1, wherein the arc is a business object.

13
14 3. The method of claim 2, further comprising:
15 at the first site, creating the public process definition;
16 distributing the public process definition to the second site;
17 at the first site, associating with the first node by shared information a first private
18 process definition containing an action preceding the transmission of the message
19 from the first site; and,
20 at the second site, associating with the second node by shared information a second
21 private process definition containing an action following the reception of the
22 message by the second site.

23
24 4. The method of claim 3, wherein the step of distributing further comprises:
25 reviewing the public process definition at the second site; and,
26 in the event of an approval of the public process definition at the second site,
27 transmitting an approval signal from the second site to the first site; and,
28 in the event of a disapproval of the public process definition at the second site,
29 transmitting a disapproval signal from the second site to the first site.

30
31 5. The method of claim 4, wherein the step of distributing further comprises:
32 in the event of receiving at the first site an approval signal from the second site,

1 transmitting a commit message to the second site;
2 installing the public process definition and the first private process definition at the first
3 site; and,
4 in the event of receiving at the second site the commit message from the first site,
5 installing the public process definition and the second private process definition
6 at the second site.

7
8 6. The method of claim 5, wherein the step of distributing further comprises:
9 in the event of receiving at the first site a disapproval signal from the second site,
10 transmitting an abort message to the second site.

11
12 7. The method of claim 3, further comprising:
13 at the first site, transforming the public process definition into a first state machine; and,
14 at the second site, transforming the public process definition into a second state machine.

15
16 8. The method of claim 3, further comprising:
17 recording in a first process execution history the execution of the first node of the public
18 process definition;
19 recording in a second process execution history the execution of the second node of the
20 public process definition;
21 auditing the first and second process execution histories.

22
23 9. A method for coordinating a process between a first site and a second site,
24 comprising:

25 creating at the first site a public process definition including a first node associated with
26 the first site, a second node associated with the second site, and an arc interposed
27 between the first node and the second node;
28 distributing the public process definition to the second site;
29 creating at the first site a first private process definition associated by shared information
30 with the first node and in which an action preceding the transmission of the
31 message from the first site is defined;
32 creating at the second site a second private process definition associated by shared

1 information with the second node and in which an action following the reception
2 of the message by the second site is defined;
3 executing the first node of the public process definition by executing the first private
4 process definition at the first site; and,
5 upon receiving the message, executing the second node of the public process definition
6 by executing the second private process definition at the second site.

7
8 10. A method for creating a process definition governing a process between a first site
9 and a second site, comprising:
10 creating at the first site a public process definition including a first node associated with
11 the first site, a second node associated with the second site, and an arc interposed
12 between the first node and the second node;
13 distributing the public process definition to the second site;
14 creating at the first site a first private process definition associated by shared information
15 with the first node and in which an action preceding the transmission of the
16 message from the first site is defined; and,
17 creating at the second site a second private process definition associated by shared
18 information with the second node and in which an action following the reception
19 of the message by the second site is defined.

1/9




FIG. 1

100

2/9

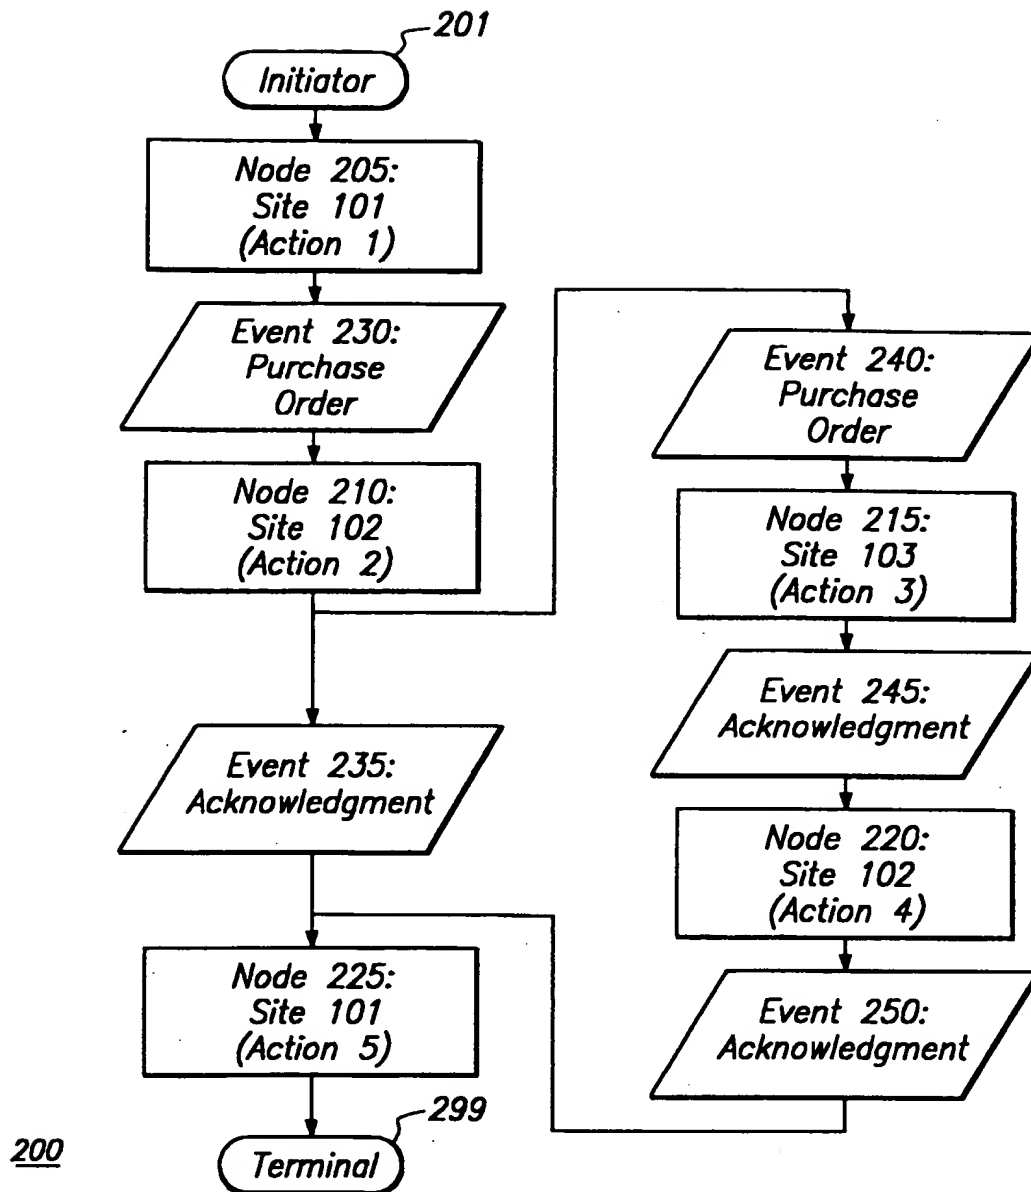


FIG. 2

3/9

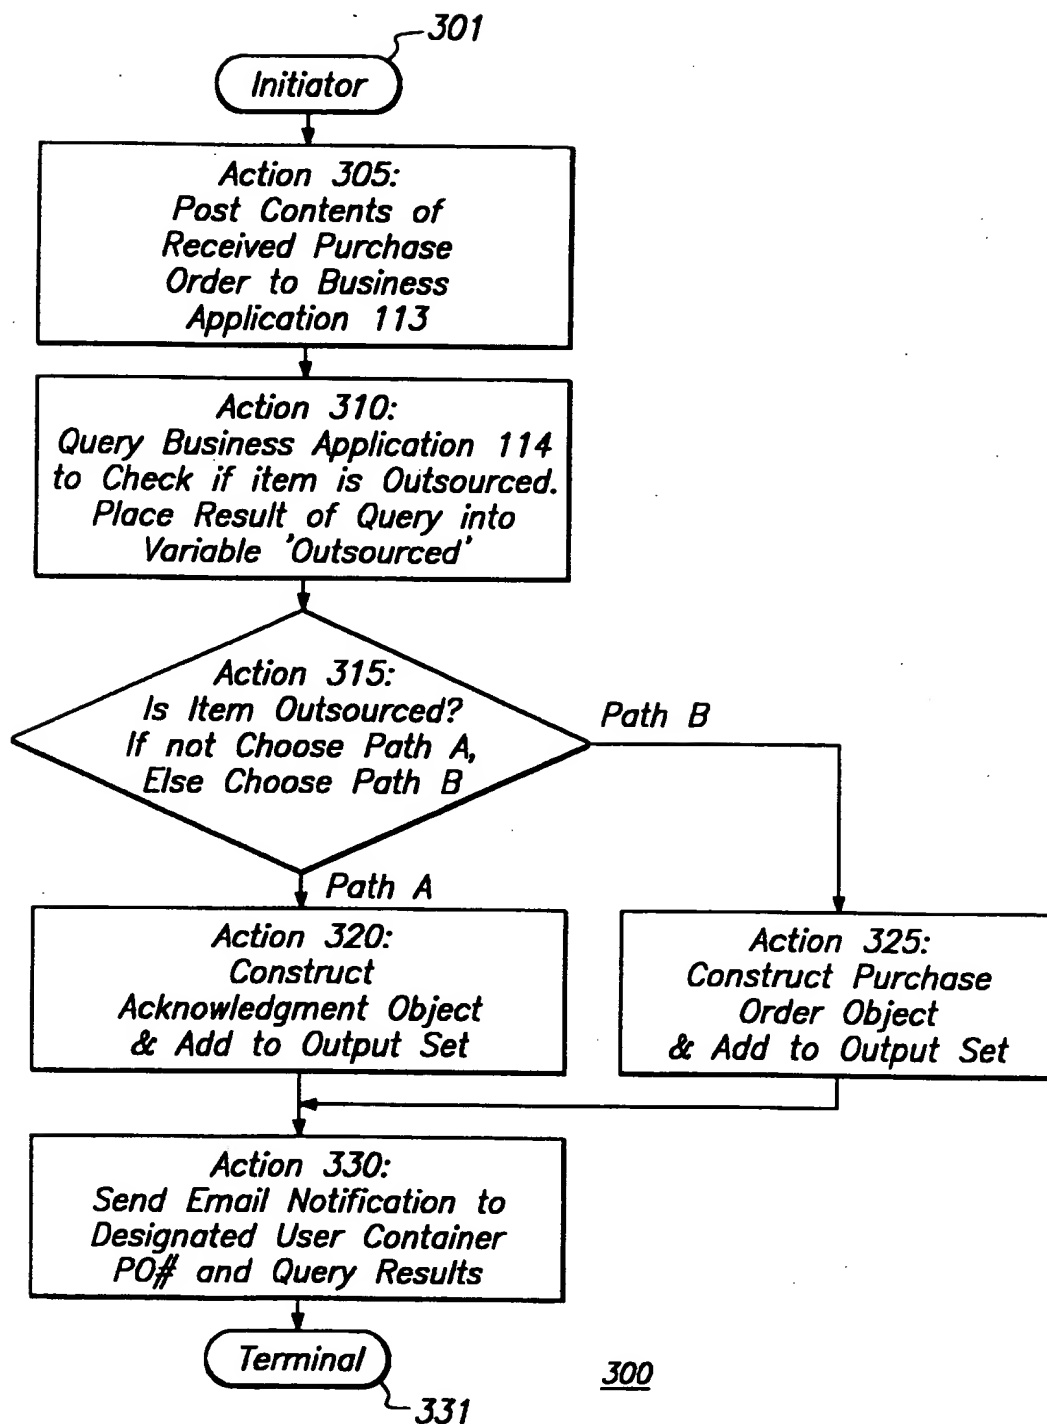



FIG. 3

4/9



5/9

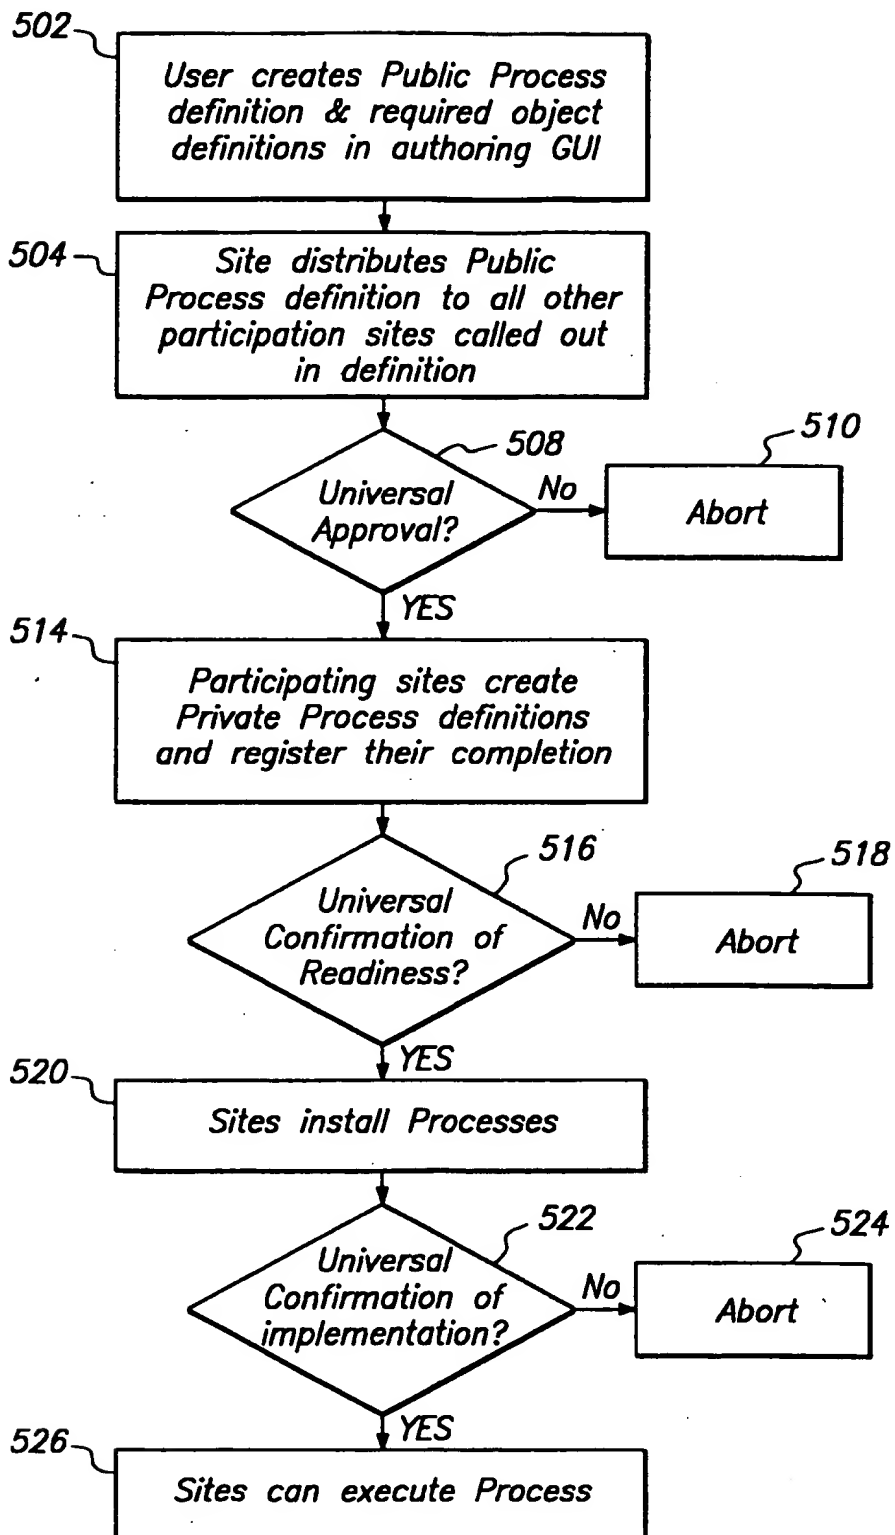
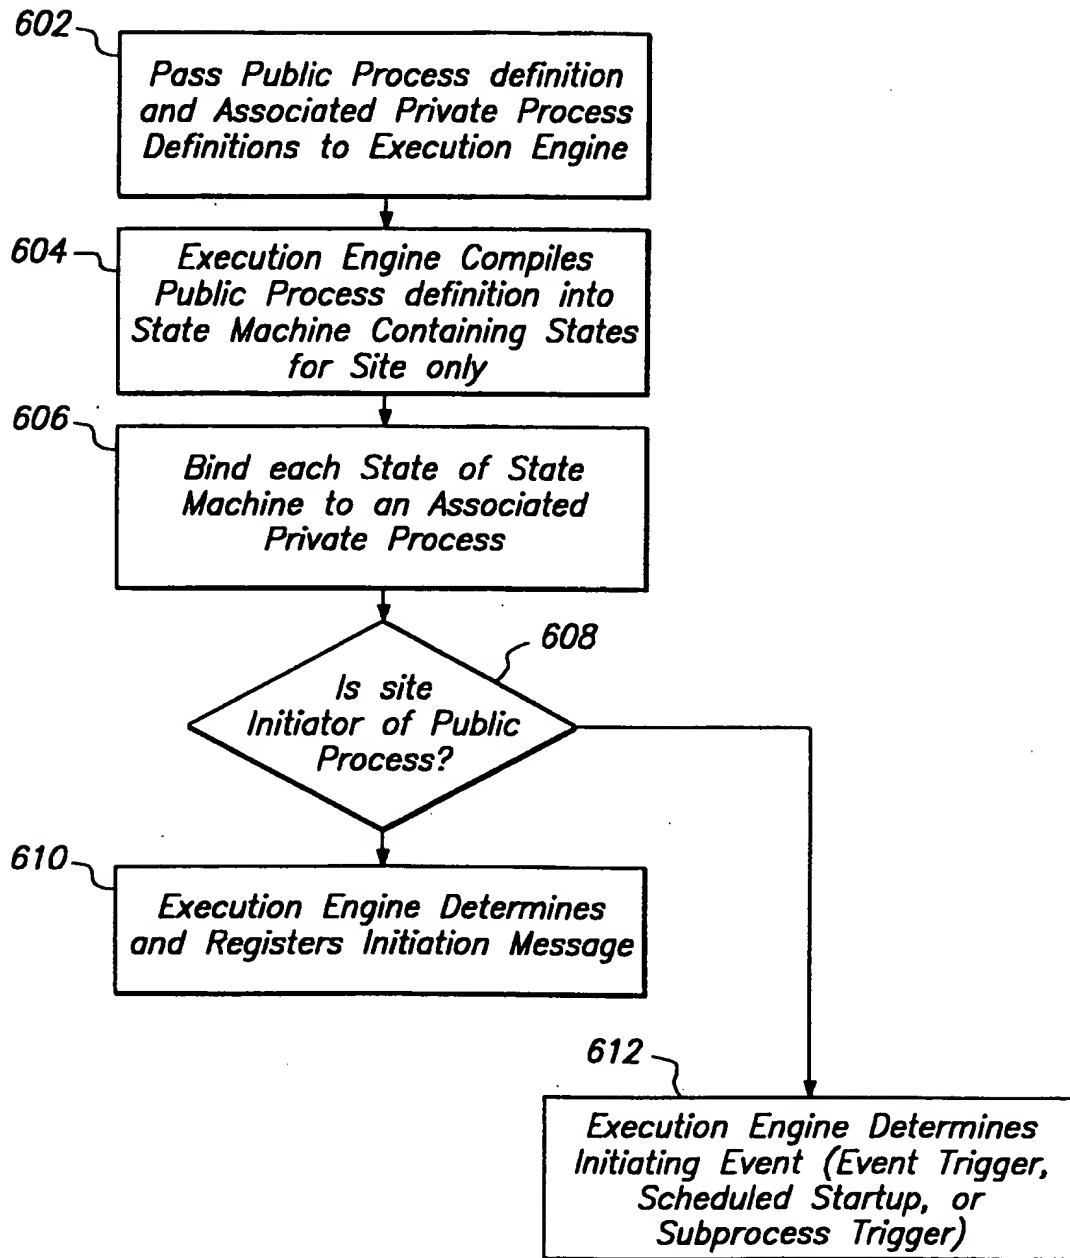


FIG. 5

6/9

**FIG. 6**

7/9

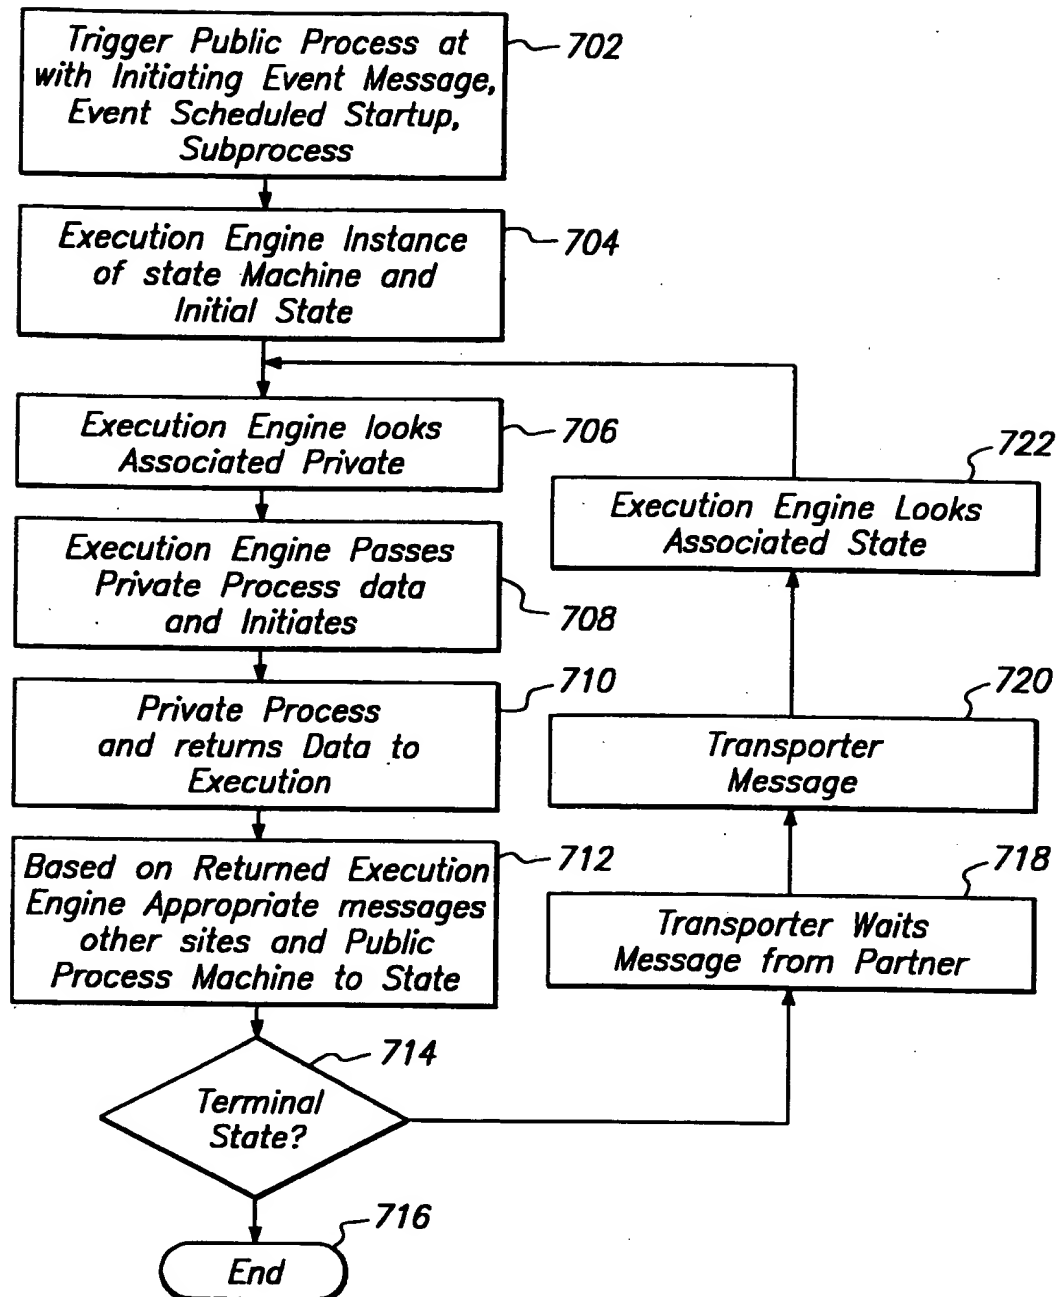
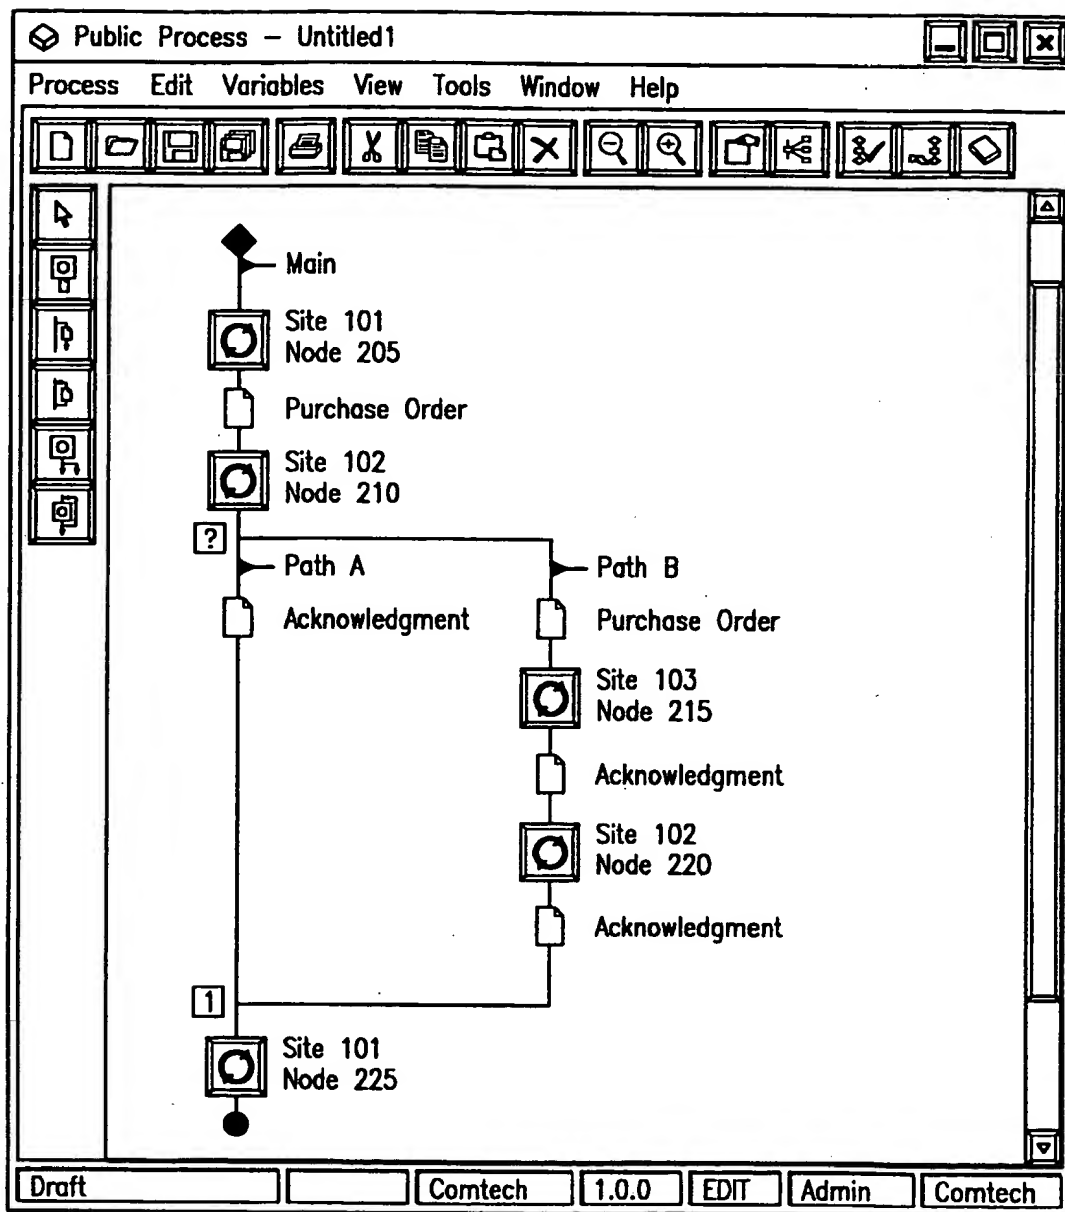
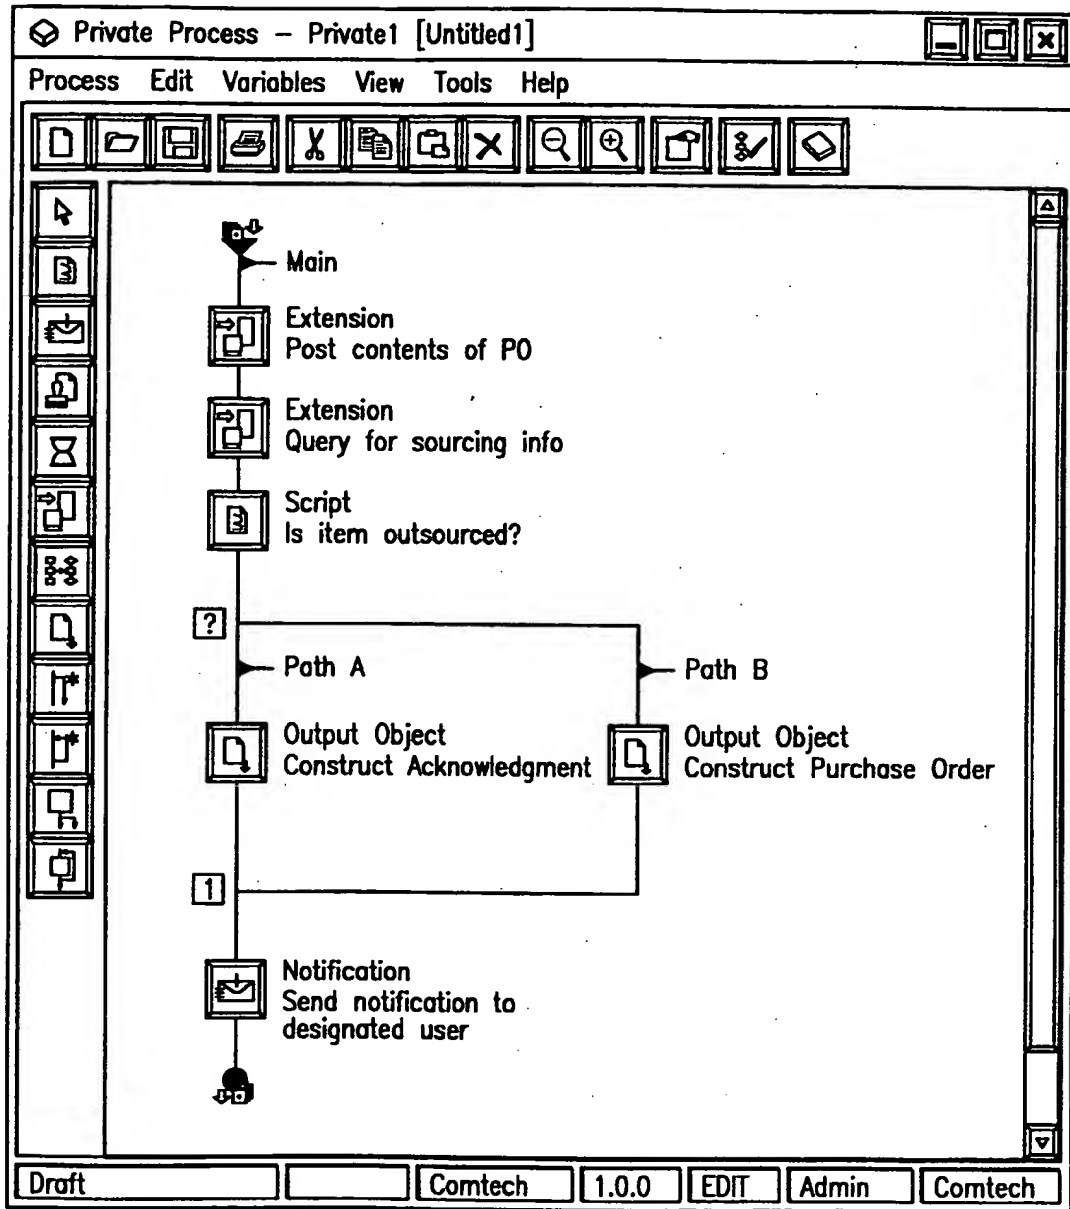


FIG. 7

8/9

**FIG. 8**

9/9

**FIG. 9**

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 98/01403

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 G06F15/16

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 390 247 A (FISCHER ADDISON M) 14 February 1995 see abstract; figures 2,20,26,27,33 see column 1, line 30 - column 3, line 25 see column 5, line 1 - column 9, line 18 see column 18, line 14 - column 29, line 53 see figures 37-40	1-10
A	DE 196 29 192 A (PREMENOS CORP) 23 January 1997 see abstract; figures 1-5,14,19,22 see figures 29,34-36 see page 2, line 11 - page 5, line 58 see page 36, line 38 - page 37, line 54 -/-	1-10



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

3 June 1998

Date of mailing of the international search report

10/06/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Gardiner, A

INTERNATIONAL SEARCH REPORT

International Application No.

PCT/US 98/01403

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>US 5 557 780 A (EDWARDS ALLAN T ET AL) 17 September 1996 see abstract; claims 6-15; figures 1-15 see column 1, line 12 - column 2, line 60 see column 3, line 16 - column 6, line 19 see column 6, line 48 - column 11, line 67</p>	1-10
A	<p>GB 2 271 002 A (DIGITAL EQUIPMENT INT) 30 March 1994 see the whole document</p>	1-10

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 98/01403

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5390247 A	14-02-1995	AU 3560793 A CA 2093094 A EP 0565314 A JP 6295286 A US 5337360 A	07-10-1993 07-10-1993 13-10-1993 21-10-1994 09-08-1994
DE 19629192 A	23-01-1997	AU 6058696 A CA 2181597 A FR 2737067 A GB 2303525 A IT MI961519 A JP 9162860 A NL 1003644 A	23-01-1997 20-01-1997 24-01-1997 19-02-1997 19-01-1998 20-06-1997 21-01-1997
US 5557780 A	17-09-1996	NONE	
GB 2271002 A	30-03-1994	NONE	